

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
14 December 2000 (14.12.2000)

PCT

(10) International Publication Number
WO 00/75962 A2

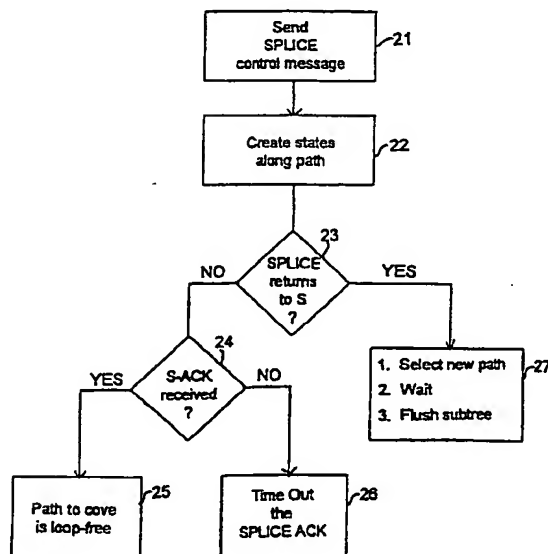
- (51) International Patent Classification⁷: **H01L** (74) Agents: HARRIS, John, D. et al.; Gowling Lafleur Henderson LLP, Suite 2600, 160 Elgin Street, Ottawa, Ontario K1P 1C3 (CA).
- (21) International Application Number: PCT/CA00/00664
- (22) International Filing Date: 7 June 2000 (07.06.2000) (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
60/137,717 7 June 1999 (07.06.1999) US (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).
- (71) Applicant: NORTEL NETWORKS LIMITED
[CA/CA]; World Trade Center of Montreal, 380 St. Antoine Street West, 8th Floor, Montreal, Quebec H2Y 3Y4 (CA).
- (72) Inventor: LEE, Cheng Yin; 7-2060 Westbury Road, Ottawa, Ontario K2C 1G8 (CA).

Published:

— Without international search report and to be republished upon receipt of that report.

[Continued on next page]

(54) Title: MECHANISM FOR SPLICING TREES



(57) Abstract: A method for detecting a routing loop when repairing/constructing a partitioned multicast distribution tree. A splice control message is sent all the way to the root-node to find out if a routing loop is formed when a node is joining the multicast distribution tree. Depending on the multicast application requirement, when a routing loop is detected during the repairing/construction of a multicast distribution tree, the severed node may abort, or the repair can be deferred until the connection to the root-node is reestablished. The loop detection and loop prevention method of the invention may be used with any protocol for constructing/splicing a partitioned tree, as a loop avoidance mechanism.

WO 00/75962 A2



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

This Page Blank (uspto)

MECHANISM FOR SPLICING TREES

Background of the Invention

5 Field of Invention

The present invention generally relates to multicast distribution trees and in particular to a mechanism for avoiding loops when a distribution tree is constructed.

Related Art

10 Simple Multicast (SM) uses a single concept working within and between domains, by offering a simplified concept for multicast distribution using the core node "C" address and the multicast address "M" ("M" is unique for a given "C"). By leaving only "C" and M" in the data-packet header, routers along a path in the multicast distribution tree have to figure out the
15 location of "C" using only "M". It reduces the states necessary in routers for supporting multicast distribution while providing a simple and reliable solution with increased speed and a lower overhead.

 A bidirectional distribution tree, as opposed to the per-source unidirectional distribution tree, is an effective solution for data delivery from
20 non-core sources. Traffic can be injected from any point and data from a source node do not have to be tunneled first to the core "C". It is also more robust since the core "C" is considered as another node in the tree and when "C" is down, the partitioned tree may still be used until the link/s to "C" is/are re-established.

25 A parent node in a multicast distribution tree transmits HEARTBEAT (HB) messages to its children at regular intervals and continues to send these messages even if it stops receiving HB messages from its parent such that a subtree may continue to function even if the core is dead. If the core is not dead, or if the path to the core has changed, the parent can simply re-join the
30 multicast tree without disrupting the nodes below. The HB message carries also a "distance from the core" indication.

 Keep-Alive (KA) messages are sent to a router by hops from further

leafs on the distribution tree, or spanning tree. The router collects KA messages from all its children ports and transmits a KA message to the router which is one hop more than the maximum hops received from any children.

5 When the distance is great in the HB and KA messages, a loop is suspected and the port removed from the tree. This is a simple loop detection mechanism and can not be used for preventing loops from forming in a spanning tree at the time the tree is repaired.

10 There are situations like unused branches on a spanning tree, loops involving member nodes, broken/changed path to core, or dead/unreachable core, which require loop detection and/or loop prevention procedures to be provided at the time a tree is repaired.

15 Methods for detecting and preventing loops formed during the construction phase of a multicast tree are known in the art. For example, US Patent No. 5,331,637 issued to the same assignee, detects transient loops in unicast routing during the construction of the multicast core based tree (CBT). Transient loops are loops that occur for a short period of time but during their existence they can damage the tree similarly to the non-transient loops. The loop prevention mechanism described in US 5,331,637 has limitations as it does not prevent loops from forming when an upstream link to the root of a tree is down, or the tree is partitioned. The entire downstream tree has to be dismantled and then the whole multicast distribution tree has to be rebuilt.

20 Suppose that node S shown in Figure 1, finds through HEARTBEAT messages that the link to the root node A (link AS) is down and tries to re-attach to the multicast distribution tree. Branch S-E-F is detached from the broken tree and the shortest path from the severed node S to A is now via node F, and other links not included in the multicast distribution tree.

25 According to the method disclosed in US 5,331,637, node S attempts to maintain its connectivity to node A and to re-attach itself to the tree by sending a graft message, or JOIN request control message. The JOIN request is received in this case by node F. It is noted that, packets flowing from node E to node F will be also forwarded back to node S since S is now a

downstream node of F.

Node F terminates the JOIN request and sends ACK-BACK to node S. A loop is now formed between nodes S and F, the control packets will be multicasted between F and S and after each looping the number of the control packets is doubled causing damage to the multicast distribution tree. Even if a data packet counter that counts down the time-to-live (TTL) at each hop is used, the number of control packets doubled at each looping is damaging the system before TTL=0.

The method in US Patent No. 5,331,637 uses the JOIN request control message which indicates that a loop exists and at the same time creates a loop as JOIN request must be terminated at node F. This method can not prevent the loop from forming when a node attempting to re-attach to the multicast distribution tree and the node where it grafts to happens to be a downstream node.

Accordingly, there is a need for a mechanism capable to avoid the formation of loops at the time a multicast distribution tree is constructed such that a node joining the tree can re-attach itself to any member node.

Summary of the Invention

The present invention seeks to alleviate totally or in part the drawbacks associated with the prior art multicast distribution tree construction.

According to one aspect of the invention, a method for avoiding routing loops when constructing/repairing a multicast distribution tree, is provided. The method comprises the steps of launching a splice message from an originating node and transmitting the splice message towards the root-node along a path, whenever the originating node intends to join said multicast distribution tree; creating forwarding states for the multicast distribution tree along the path; determining if the splice message is returned to the originating node; declaring loop-free if the splice message is not returned to the originating node and a splice acknowledgment message generated by the root-node is received by the originating node; and terminate

the constructing/repairing the multicast distribution tree if the splice message is received back by the originating node.

The present is not limited to the features disclosed in the "Summary of the Invention" section; it nonetheless may reside on sub-combinations of the disclosed features.

Description of the Drawings

The present invention and its advantages will be better understood from the following description of the embodiments of the invention illustrated in the drawings, where:

Figure 1 is an illustration of a branch of a multicast distribution tree; and

Figure 2 is a flow chart illustrating the method of avoiding loops during the construction of a multicast distribution tree.

Similar references are used in different figures to denote similar components.

Detailed Description of the Preferred Embodiments

The following description relates to preferred embodiments of the invention by way of example only and without limitation to the combination of features necessary for carrying the invention into effect.

When a multicast distribution tree is partitioned, the available routing protocol will try to repair the tree with as little as possible disruption of the traffic to the nodes below. To repair such a partitioned multicast tree the downstream subtree, or the node severed from the main multicast distribution tree, has to be re-attached to the main tree provided a loop is not formed in the process.

According to the invention, a method for detecting and avoiding loops from forming when two multicast distribution trees are spliced, is provided. Depending on the multicast application requirement, when a routing loop is detected during the repairing of a multicast tree, the severed node may abort,

or the repair can be deferred until the connection to the root-node is re-established.

A SPLICE control message is used as a "loop probe" all the way to the root to discover whether or not there is a potential routing loop forming
5 between a severed node S and root-node A of Figure 1.

The SPLICE message, unlike the JOIN request, is not terminated by the intermediate routers which are members of the multicast tree, as they already have forwarding states for the tree. Thus, SPLICE can be sent all the way to the root to make sure that the path to the root is loop-free.

10 Node S, see Figure 1, discovers the tree is partitioned, eg. link AS is down and sends SPLICE towards root-node R. In order to re-attach to the multicast distribution tree, the severed node S launches the SPLICE message, step 21 of Figure 2, to detect if the path from the severed node S to the root-node R is loop-free. The SPLICE message is forwarded all the way
15 to the core creating forwarding states for the multicast tree at nodes along its path to the root-node R, step 22.

At step 23, node S joining the multicast distribution tree determines whether SPLICE returns to S. If SPLICE hits a downstream node, the SPLICE message returns to S.

20 If SPLICE reaches the root-node R, node R creates forwarding states (as if it has received a JOIN-REQUEST message) and sends a SPLICE acknowledge message to S, step 24, to acknowledge the receipt of the SPLICE message.

25 If the severed node S receives a SPLICE ACK message, step 25, this indicates the path to the root is loop-free and the joining node may rejoin the multicast distribution tree using the conventional JOIN request control message to splice the subtree. In other words, node S has learned that it can reattach to the multicast tree without causing loops. As well, node S makes the transient forwarding states (created earlier) permanent.

30 If the path to the root of the multicast distribution tree loops, the severed node S will not receive any SPLICE ACK message, step 26, and

should time out the waiting for the SPLICE ACK message.

At step 22, SPLICE message is forwarded all the way to the core, creating state where there is no existing multicast state, but merely forwarded to the root when there is an existing multicast state. Advantageously, the method of the invention allows the SPLICE message to create transient forwarding states (like the JOIN request message), but forward the SPLICE message towards the root even if the router receiving the SPLICE message has an existing forwarding state for the multicast group.

For example, when node G receives the SPLICE message, creates transient forwarding states (as if it has received a JOIN-REQUEST message) and forwards the SPLICE message to towards the root R. As well, node G makes the transient forwarding state permanent (as if it has received a JOIN-ACK message) and forwards the SPLICE ACK to S.

As mentioned before, if the SPLICE message hits a downstream node, the message is forwarded until it reaches the router which originated the SPLICE message. The originating node has now learned that it can not splice the subtree without causing loops, step 27.

For example, node F receives the SPLICE message and forwards the SPLICE message without creating forwarding states, towards R via E. Node also E forwards the SPLICE message without creating forwarding states, towards R. Node S receives the SPLICE message and finds out that it is the originator of this message. Node C has learned that it can not use this path to reattach to the multicast tree as it creates a routing loop.

Depending on application requirement conveyed to nodes from root via heartbeat messages, the multicast tree construction protocol can be configured to:

- (1) Attempt to rejoin the multicast tree via another next hop (e.g equal cost route to R if available); or
- (2) Wait for a pre-defined period before attempting to rejoin R via the same next hop or another next hop, allowing the subtree to continue to function separately, but attempting to SPLICE the tree again when the

unicast route to core changes, or is reestablished in the case where there is no route available to the root; or

(3) FLUSH the sub-tree with parent node S and let the receivers downstream of S to individually rejoin the multicast tree

5 The SPLICE messages are encapsulated in an IP header. The protocol ID of the IP header is set to the tree construction protocol using this splice mechanism.

 The fields required in a SPLICE message are:

10 - SPLICE Message: This contains the SPLICE Message type, the message length, address of the router which originates the SPLICE message, the multicast group address.

 The IP source address is set to the node originating the message and the IP destination address is set to the root.

15 - SPLICE Acknowledgment Message This contains the SPLICE Acknowledgment Message type, the message length, the address of the router which originates the SPLICE Acknowledge message, multicast group address. The IP source address is set to the root node sending the SPLICE Acknowledge message and the IP destination address is set to the node which originates the SPLICE message.

20 A method for detecting a routing loop when repairing/constructing a partitioned multicast distribution tree, was presented. A splice control message is sent all the way to the root-node to find out if a routing loop is formed when a node is joining the multicast distribution tree. Depending on the multicast application requirement, when a routing loop is detected during
25 the repairing/construction of a multicast distribution tree, the severed node may abort, or the repair can be deferred until the connection to the root-node is re-established. The loop detection and loop prevention method of the invention may be used with any protocol for constructing/splicing a partitioned tree, as a loop avoidance mechanism.

30 In this way, the control message used to repair a partitioned multicast tree message is not overloaded to function as both a graft message and a

loop detection message. Without the capability of detecting a routing loop during the repair of a multicast distribution tree, the conventional multicast distribution tree construction protocols have to remove the downstream subtree entirely.

5 The invention can be implemented in digital electronic circuitry, or in computer hardware, firmware, software, or in combinations thereof.

Apparatus of the invention can be implemented in a computer program product tangibly embodied in a machine-readable storage device for execution by a programmable processor; and the actions can be performed
10 by a programmable processor executing a program of instructions by operating on input data and generating output.

 The invention can be implemented advantageously in one or more computer programs that are executable on a programmable system including at least one programmable processor coupled to receive data and instructions
15 from, and to transmit data and instructions to, a data storage system, at least one input device, and at least one output device. Each program can be implemented in a high-level procedural or object oriented programming language, or in assembly or machine language if desired; and in any case, the language can be a compiled or interpreted language.

20 Suitable processors include, by way of example, both general and special purpose microprocessors. Generally, a processor will receive instructions and data from a read-only memory and/or a random access memory.

 Generally, the system will include one or more mass storage devices
25 suitable for tangibly embodying computer program instructions and data include all forms of non-volatile memory, including by way of example semiconductor memory devices, such as EPROM, EEPROM, and flash memory devices; magnetic disks such as internal hard disks and removable disks; magneto-optical disks; and CD-ROM disks. Any of the foregoing can be
30 supplemented by, or incorporated in, ASICs (application-specific integrated circuits).

Numerous modifications, adaptations, and variations may be made to the particular embodiments of the invention without departing from the scope of the invention which is defined in the claims.

5

10

15

20

CLAIMS:

1. A method for avoiding routing loops when constructing/repairing a multicast distribution tree, comprising the steps of:

a) launching a splice message from an originating node and transmitting said splice message towards the root-node along a path, whenever said originating node intends to join said multicast distribution tree;

b) creating forwarding states for said multicast distribution tree along said path;

c) determining if said splice message is returned to said originating node;

d) declaring loop-free if said splice message is not returned to said originating node and a splice acknowledgment message generated by said root-node is received by said originating node; and

e) terminate said constructing/repairing said multicast distribution tree if said splice message is received back by said originating node.

2. The method of claim 1, wherein said step of creating forwarding states along said path including:

- creating said forwarding states at nodes where said forwarding states are not existing,
- forwarding said splice message to said root-node, and
- forwarding said splice acknowledgment back to said originating node.

3. The method of claim 1, wherein said step of creating forwarding states along said path including:

- creating transient forwarding states at nodes where said forwarding states

are existing,

- forwarding said splice message to said root-node, and
- forwarding said splice acknowledgment back to said originating node.

5 4. The method of claim 3, further comprising making said transient state permanent after performing step (d).

5. The method of claim 1, wherein said step (e) further comprising attempting to join said multicast distribution tree via another hop.

10

6. The method of claim 1, wherein said step (e) further comprising waiting before attempting to join said multicast distribution tree via same hop.

15

7. The method of claim 1, wherein said step (e) further comprising flushing the subtree with said originating node as the parent node and allowing downstream nodes to individually join said multicast distribution tree.

8. A computer-readable medium containing computer executable instructions for performing the steps of:

20

- launching a splice message from an originating node and transmitting said splice message towards the root-node along a path, whenever said originating node intends to join said multicast distribution tree;

25

- creating forwarding states for said multicast distribution tree along said path;

- determining if said splice message is returned to said originating node;
- declaring loop-free if said splice message is not returned to said originating node and a splice acknowledgment message generated by said root-node is received by said originating node; and

30

- terminate said constructing/repairing said multicast distribution tree if said splice message is received back by said originating node.




FIG. 1

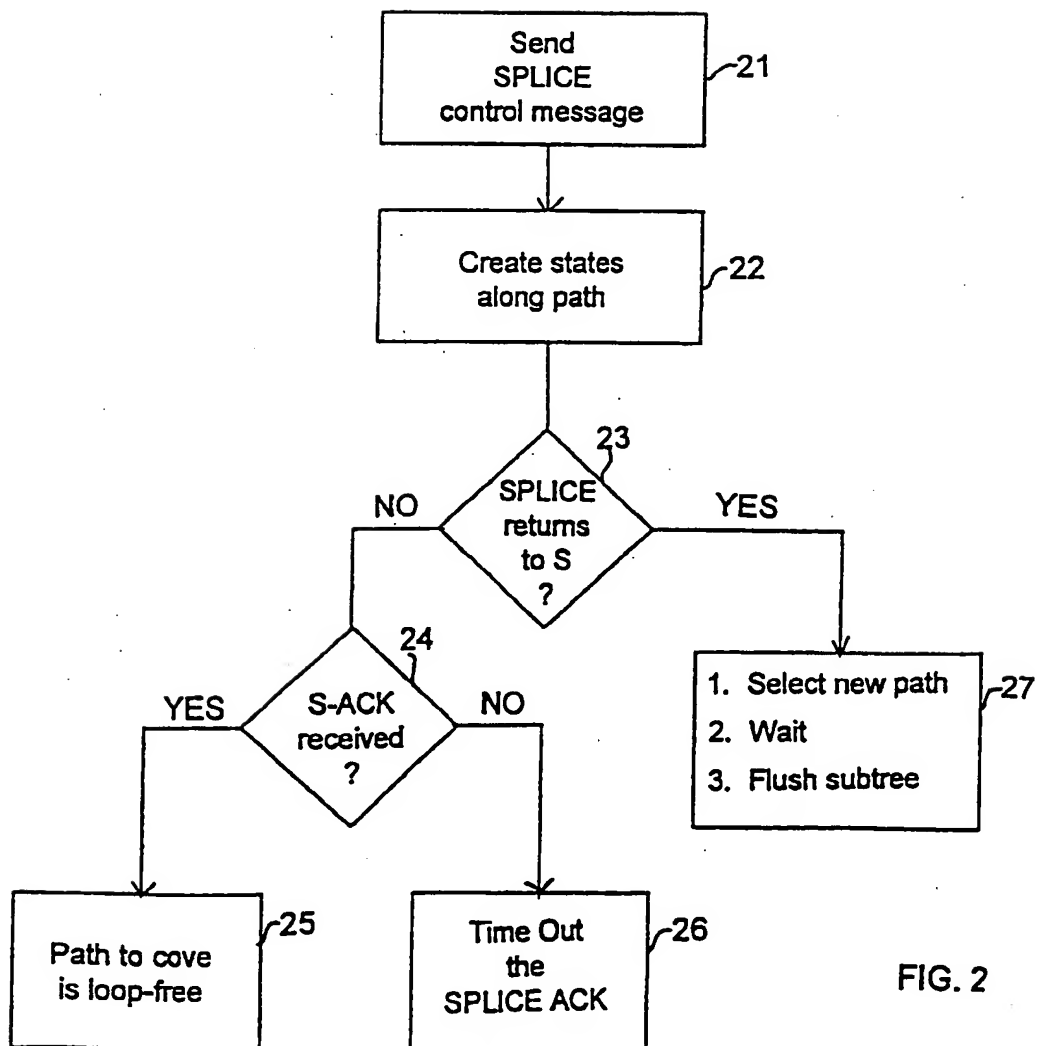


FIG. 2